NCC98 Network Backbone Design

Operations Concept

Updated 12/1/97

Introduction

- ❑ Purpose: This document will provide an overview of the operational concepts involving the NCC98 network backbone equipment such as the high-end switches and associated ancillary equipment.

- ❑ Scope: The operational concepts described in this document are limited to the NCC98 Network Backbone only (i.e...., LAN infrastructure such as switches, hubs, LAN interfaces on critical systems, cabling).

Current NCC Network Backbone Overview

- The current backbone topology is based on physically isolated shared 10 MHz Ethernet segments to which individual nodes are accessible via a cross-point matrix switch (NACC) which is manually operated and controlled. A node can be connected to only one Ethernet segment (LAN) at a time. The connection stays in place until manually removed by an operator.
- Switching, as defined by this framework, is the result of manually (via a console or hand-patch) connecting a node into a predefined segment or LAN such that the node now shares the segment's 10 MHz of bandwidth. This means that the node can listen to every other node on that segment as well as talk to every other node on that segment.
- Failover of a redundant critical system is most often a manual process which requires manual detection, notification, and operation whether performed on a console or hand patched.
- Nodal access via the backbone to all portions of the NCC (ENCC, DT&T, OPS) exists via AIS level changes if necessary; e.g. - you can get to CCS 3 in the DT&T if necessary from the OCR (OPS) by instituting the proper procedures for security level change and the proper reconfiguration being made in the NACC switch and on CCS 3.

Current NCC Network Backbone Overview (cont.)

- ❑ The NACC switch was developed explicitly for the Red/Black era incorporating shielded layers within PCB cards which carried traffic from separate isolated LANs in addition to the shielding the chassis itself provides. With the advent of de-classification, these features of traffic isolation are no longer a foremost requirement to be met by the network.
- ❑ Bandwidth usage and LAN performance within the Operational LANs have not been a noteworthy issue under the current topology. Bandwidth usage has been consistently minimal with measurements typically averaging less than or equal to 10%.
- ❑ Backups (Operational LANs) are performed locally by each system with no backup traffic flowing over the LAN. System backups are a manual process for each individual system/machine which is a time consuming administrative process. System information/data is stored on tape which requires space for storage and is physically dispersed in various areas.
- ❑ Network management of the LAN backbone is limited to the NACC console GUI which gives the operator link status and sufficient data to determine which LAN a node is connected to. In addition, each link has a packet count indicator. Alarms and logging features are also supported.

NCC98 Network Backbone Overview

- ❑ **In order to choose an appropriate LAN backbone design and underlying backbone technology, issues such as user requirements, past versus predicted performance statistics and analysis, and nodal usage capabilities needed to be evaluated. From this evaluation, the following assumptions could be made which will help guide the backbone design and subsequent choice of LAN technology best suited for application.**

- ❑ **Assumptions:**
  - » All critical components or systems will be redundant in Operations:
    - – High-end LAN switches
    - – K Servers (clustered to support NCCDS DBase, SPSR, NSM, et.al.)
    - – NPG
    - – Firewall
    - – TUT, FTP Server
    - – CCS (VAX 8850)
    - – "SCD" (NFE replacement version)
    - – HP Workstations

NCC98 Network Backbone Overview (cont.)

❑ Assumptions (continued)

» ANCC will not necessarily maintain redundant critical systems

» Automatic failover for all critical systems

» High Availability (HA) software will be used for the following critical systems:

- K Server Cluster will use HP's ServiceGuard
- NPG, and Firewall using custom software for HA

» Nascom IONET connections for NCC98 will be running a dynamic routing protocol for automatic failover to the redundant circuit (includes firewall, link, and router)

» Redundant gateway (firewall) is powered up but cannot be active (due to duplicate message transmission)...this requires use of high availability software on the firewall

» Expansion rate projections (per port): 30% for Ethernet/Fast Ethernet over proposed design for technology life cycle (5-7 years)

NCC98 Network Backbone Overview (cont.)

❑ Assumptions (continued)

- » Current node count projection for NCC98 Operations (port requirements with expansion included)
  - – Ethernet: 42 ports total or 21 ports per switch
  - – Fast Ethernet: 16 ports total or 8 ports per switch
- » Performance projections by the NSM group anticipate the potential for heavy traffic volume during backups of critical systems/applications
  - – Shared Ethernet hubs will be insufficient as they are not scaleable, are not designed for high availability, and have limited additional features
  - – Switched LAN technology is the recommended choice
    - Capable of high performance (a switched connection with wire speed throughput)
    - 5-10 year technology life cycles
    - Versatile support for Ethernet, Fast Ethernet, Gigabit Ethernet, Token Ring, ATM, FDDI/CDDI, or WAN connectivity
    - Highly fault tolerant and fully modular (e.g., distributed mgmt, redt. power supp.)
    - SNMP manageable
    - Highly functional (VLANs, layer 3 switching, intelligent, connection oriented)

NCC98 Network Backbone Design Project

Proposed NCC98 Backbone Design

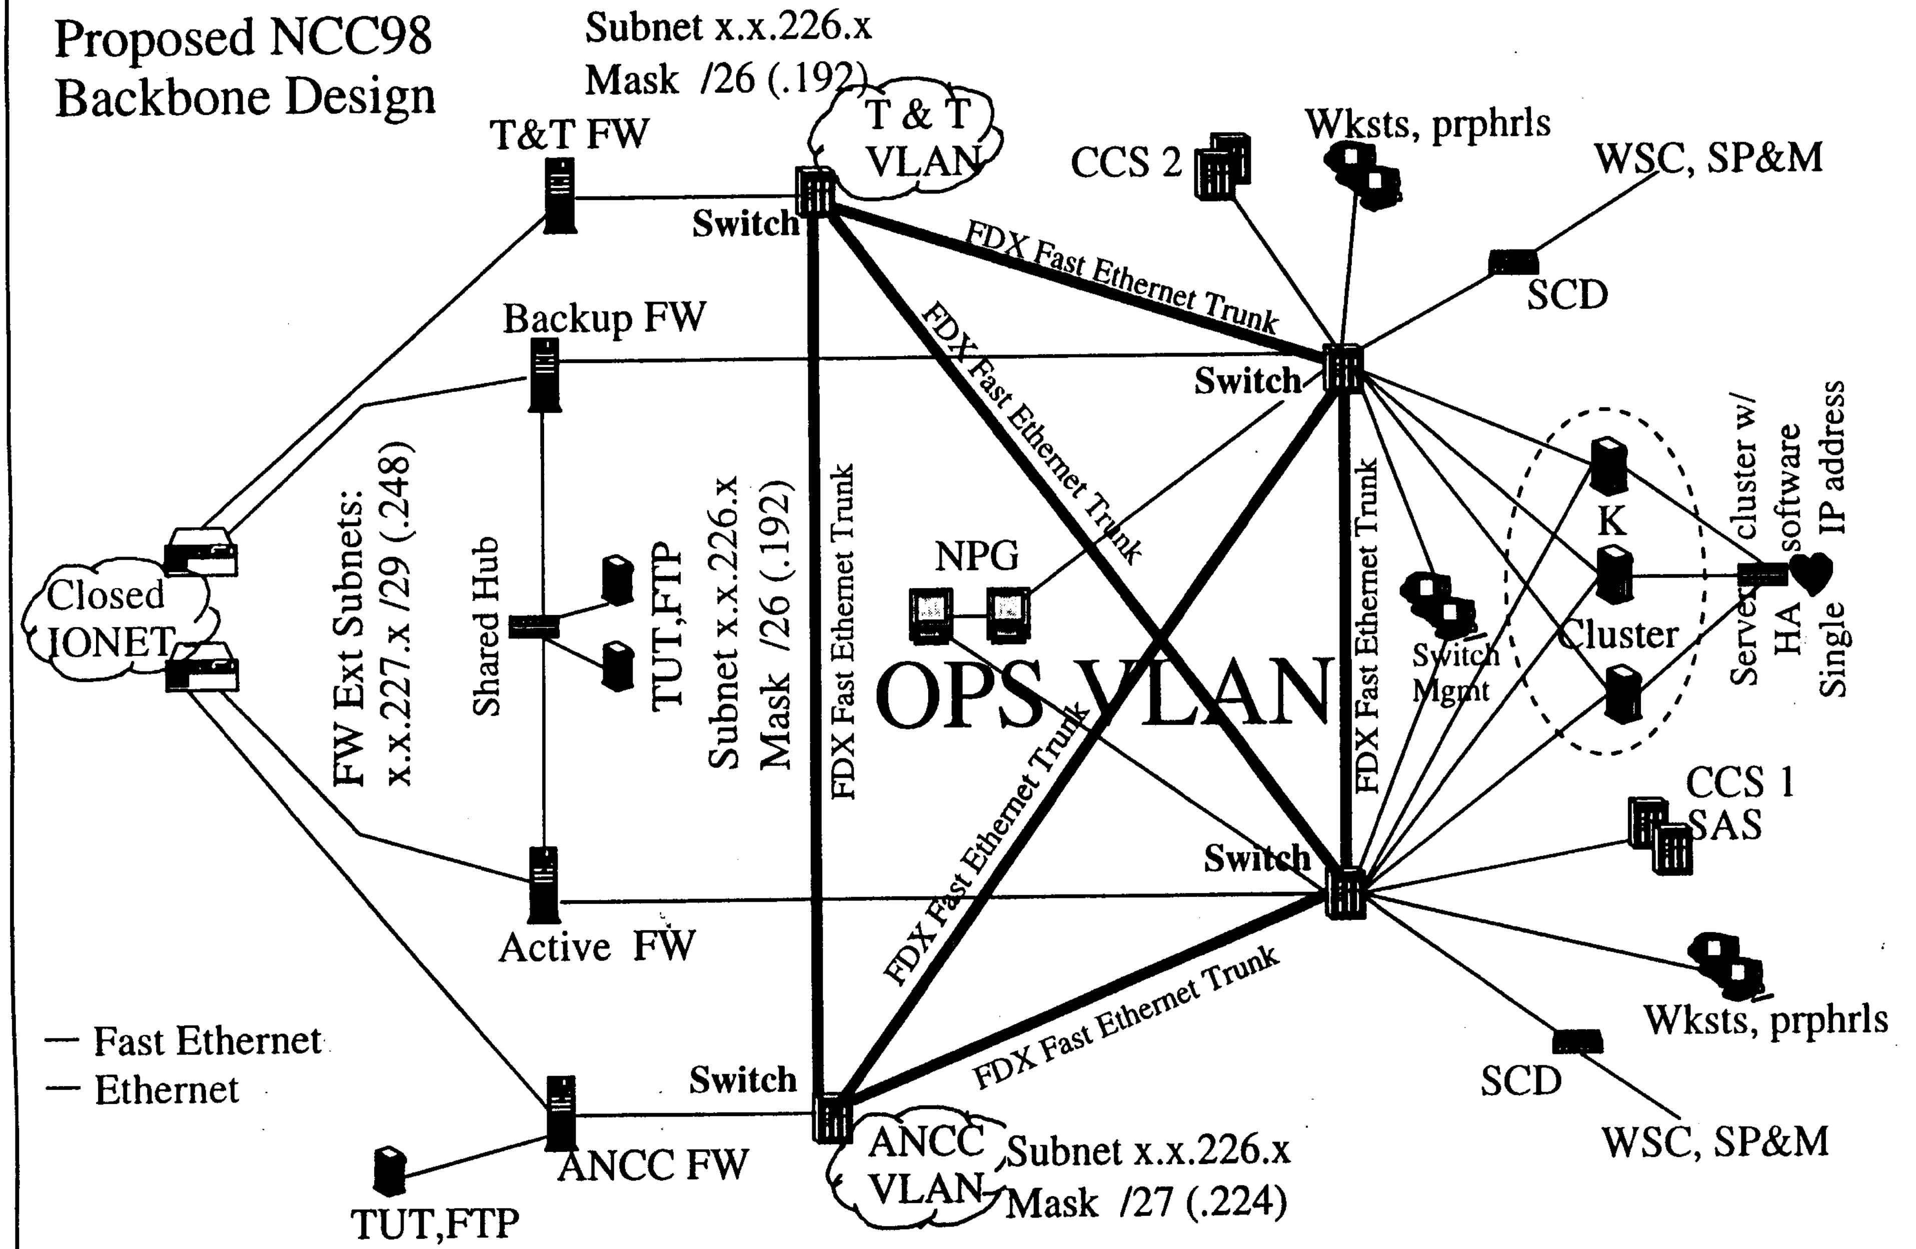

Topology Highlights

❑ Diversely Run, Redundant IONET Connections

- » Connections to the external IP world for NCC98 will be via the Closed IONET as provided by Nascom. These firewall protected NCC connections are actually part of the IP Transition network which is still under development. At the present time, there will be four distinct circuits as depicted in the diagram on the preceding page. One each to the Test & Training (T&T) subnet and Auxiliary NCC (ANCC) subnet, and two to the Operational (OPS) subnet.
- » The two circuits running to the OPS subnet are diversely run (physically separate paths to Nascom) and redundant such that one circuit will be a dynamic backup to the other. Additionally, they each run to separate routers in Nascom. The use of link-state routing protocols in conjunction with high availability software will enable these two links to function as a fully redundant, automatic failover system for external connectivity to/from the NCC.
- » Note that the existing serial communications based circuits which support the White Sands Complex (WSC) and Special Projects & Missions (SP&M) will still exist and will continue to be fed into NCC98 via conversion devices.

Topology Highlights (cont.)

❑ Security Changes Provide Increased Flexibility

» As presented at the *Equipment Suite & Networks Operational Concepts Review* back in August, the separation of the Development environment from the Test & Training environment enables the NCC to have more flexibility in the use of its resources. By qualifying the T&T as an AIS level 3 environment, OPS will be able to utilize more easily the resources within the T&T and ANCC as required. This will be possible by the use of some of the state of the art features provided by the backbone's new workhorses, high-end LAN switches. These high-end LAN switches will be configured as a fully meshed switch topology thus providing multiple active, load sharing/balancing trunks between each of the four switches. Although the load sharing/balancing aspect will not be of much impact to our design (since it isn't expected that there will be large volumes of data transfer between T&T and OPS or between ANCC and OPS), it is the fact that multiple active trunk links provide a high degree of reliability for the design overall.

Topology Highlights (cont.)

❑ High-end LAN switches

» Today's bandwidth hungry client/server applications running on larger and more powerful endstations with larger bus structures will quickly fill a shared bandwidth LAN when enough are operating at the same time. Throw a few backups on the same LAN and you have the necessary ingredients for network congestion. Switched LAN technology is the logical step in the migration from a shared LAN technology. By breaking up the collision domain with dynamically assigned connections, a switch provides the necessary aggregate throughput to reduce network congestion. However, there is still the problem with many connections being made with a single node, a server for example. This problem is solved with the addition of a larger amount of switched bandwidth to the popular node, for example Fast Ethernet, FDDI, or ATM.

» Based on the scenario above, a typical switched connection using Ethernet for the majority of nodes and Fast Ethernet on the most popular nodes will be dynamically set up by the switch via a layer 2 (MAC address) table lookup, dedicate 10 Mbps of bandwidth to the connection, and once complete, "tear down" the connection (i.e., it's made virtually as opposed to physical dedication of bus wires to each and every port).

Topology Highlights (cont.)

» High-end switch designs today go far beyond simply providing network collision domain segmentation (bridging) as when they were originally introduced. High-end switches provide a wide range of features and functions that can be used to make a network much more robust, resilient, and efficient. High-end switches are typically based on a modular chassis design which supports a combination of Ethernet, Fast Ethernet, Gigabit Ethernet, Token Ring, ATM, WAN, and even some other types of LAN technologies thus providing a high degree of scalability in today's ever changing network centric world. The switching fabric is typically ASIC based which offers superior performance over software based fabrics. Most of these switches are highly redundant, offering dual redundant power supplies, dual management or supervisor engines, and dual AC power feeds. High reliability comes in the form of hot swappable modules and MTBF figures typically in the 10,000's and 100,000's of hours.

» Features such as virtual LAN (VLAN) support, provisions for individual node network analysis, SNMP/RMON supported network management, network management via secure means (serial, out of band, or isolated backplane), layer 3 switching (protocol-less routing with the performance of a switch), and even full routing capabilities all combine to make today's

Topology Highlights (cont.)

high-end switches extremely versatile, high performance networking components.

» As depicted by the Backbone diagram on page 8, NCC98 uses four high-end LAN switches employing both switched Ethernet and switched Fast Ethernet. Multiple Fast Ethernet trunks between switches form the backbone of the entire network providing multiple links between switches. The recommended switch vendor's switch software runs an OSPF based link state protocol over these trunks such that both traffic load sharing and resiliency are satisfied as fundamental design criteria. Fast Ethernet links will also be used for the clustered K servers which will have dual Fast Ethernet NICs with "mean time before failure" (MTBF) figures measuring 166,000 hours calculated (source: Hewlett Packard). Fast Ethernet expandability will be provided not only in resident modules, but with extra slots for additional switch modules if necessary. All other nodes including external links and workstations will be switched 10MHz Ethernet interfaces.

» The Fast Ethernet trunks and server links have the capability to be run in a full-duplex mode which can effectively boost bandwidth to 200 MHz if congestion proves to be a measurable inefficiency.

Topology Highlights (cont.)

» All of the high-end LAN switches which were evaluated by this project are modular in design (although to differing extents). Other sought after hardware criteria were dual redundant power supplies with separate AC power feeds, dual supervisor/management engines or distributed management, absence of or minimal logic built into the chassis, high aggregate bandwidth backplane, ASIC based switching fabric, scalability for LAN technology and port density, and optional management connections. The switch chosen for recommendation to the Project Office for implementation is equal to or better than the competition in all of these categories.

❑ Virtual LAN (VLAN) technology

» Virtual LAN or VLAN technology is a feature added to LAN switches which gives the network administrator the ability to logically group nodes such that their effective domain becomes inclusive to only nodes assigned to the VLAN. Rather than group nodes physically, which requires labor intensive resources, nodes can be grouped "virtually" or logically within the switch via a number of different parameters such as by the actual port they're attached to on the switch, or by their MAC address, or by their IP address, or even by the protocol or application they're running.

Topology Highlights (cont.)

» VLANs form logical broadcast domains (layer 2 implied) such that nodes within a specific VLAN can only connect to other nodes in the same VLAN. Thus they form boundaries of traffic isolation. Connectivity can be established outside of the VLAN by use of a routing protocol (layer 3 implied). This makes the use of VLANs an efficient administrative tool due to the ease of logically grouping nodes as opposed to physically grouping nodes. A benefit of forming smaller broadcast domains is that it makes LAN bandwidth more efficient, resulting in less collisions within a domain. Another benefit is that it gives the administrator the ability to control LAN security via the management workstation which is running the LAN equipment vendor's VLAN management software.

» Network or LAN equipment vendors implement VLANs with various features and functions. As a result, many are vendor specific or proprietary and will only work with their own equipment. Within the last year however, standardization efforts have resulted in a few vendors with the capability to manage other vendors VLANs as long as they too follow the specification (IEEE 802.3Q). Although more often than not, it is the features and functions which are beyond the scope of the spec which can make a particular vendor's VLAN flavor more desirable for a particular

Topology Highlights (cont.)

application or design. This is the case with NCC98.

» It is intended that NCC98 will use VLAN technology for all of it's primary functions such as administrative efficiency, traffic segmentation, and security in addition to many vendor specific features such as call tapping, multiple VLAN per port support, true layer 3 switching capabilities/performance, ability to quickly call up predefined VLAN configurations, and connection-oriented traffic flows.

» NCC98 will use three VLANs based on IP subnet. One will be designated for the Test & Training (T&T) environment, one for the Auxiliary NCC (ANCC), and one for the Operational environment (OPS). Via the switch vendor's VLAN management software, NSM operators will have the capability to make available a resource within either the T&T or ANCC environment such that communications will be established with that node and the rest of the OPS VLAN while severing communications with the chosen nodes former VLAN members. This action will be performed from the VLAN management software's GUI application as a simple "drag and drop" procedure. Although to fully complete the action, a simple three to five step procedure may have to be followed which will be outlined in the forthcoming *NCC98 Backbone/Switch Management User's Guide.*

Topology Highlights (cont.)

- » If necessary, the recommended switch manufacturer for NCC98 offers multiple VLAN per port capability such that the resource could still be used by two or more VLANs while still preserving the intended VLAN traffic isolation capabilities.
- » It should be noted that the security benefits of VLAN technology should be used to the fullest extent in NCC98. Additional VLAN(s) for network management/administrative functions could be set up to isolate this traffic from operational traffic thus providing an increased layer of internal security. This could be accomplished with the recommended switch vendor's VLAN technology only as it provides the capability to handle multiple VLANs per port.
- » Although it is not this project's responsibility to define network management issues, it is recommended that switch/VLAN management capabilities be provided in each of the three environments (T&T, OPS, ANCC) with overall control of all three environments run out of OPS (with failover capability to the ANCC). This would allow management of resources within the T&T to be handled internally as testing needs arose but restrict any management capability outside of their subnet/VLAN via security protocols (e.g., password protection, other).

Topology Highlights (cont.)

» The concept of switched network analysis is almost an oxymoron. The problem is that network analysis works great with a shared bandwidth LAN technology because all of the traffic on the LAN can be seen by any node attached to the LAN. On a switch however, traffic is isolated due to the nature of switched connections which effectively form dynamic collision domain segments for the duration of the connection. This makes it very tough to analyze traffic on the switch. Most of the high-end switch vendors support network/traffic analysis through the use of some sort traffic mirroring or copying feature which will copy the traffic on a given port or connection to another port where the network analyzer should be located. This is accomplished through the switch vendor's management software. Some switch vendor's can even monitor multiple ports or connections at the same time which closer emulates the shared LAN capability.

» The recommended switch vendor's VLAN management software provides network analysis capability through a feature known as "Call Tapping". This is a term typically associated with connection oriented communications such as ATM and frame relay and is likewise appropriate for this vendor's connection oriented switching capabilities.

Topology Highlights (cont.)

□ Support of Heartbeat for the K Cluster

» The K-server cluster requires dual heartbeat networks in order to sustain high availability switchover capabilities. One of the heartbeat networks will be physically cabled to a SNMP manageable shared LAN hub. Each of the three K servers will have a physical NIC (network interface card) dedicated to this heartbeat LAN connection. The second heartbeat network will utilize the switched data network (on which operational traffic flows). Again, the multiple VLAN per port feature of the recommended switch vendor's VLAN technology is appropriate for use here, however, this does not guarantee that potential congestion problems wouldn't interfere with the arrival/delivery of a heartbeat "pulse" (packet). Only a physically distinct NIC or prioritization of traffic (Resource Reservation Protocol aka, RSVP) would completely guarantee delivery.

□ Additional Heartbeat Connections for HA Systems

» Additional heartbeat connections for the other HA systems such as the NPG and the Firewall will be dual point to point twisted pair cable (null/crossover wiring) connections. Each will run directly from a dedicated NIC on machine A to a dedicated NIC on machine B.